Productivity and Quality in Health Care: Evidence from the Dialysis Industry*

Paul L.E. Grieco† Ryan C. McDevitt‡

May 2012

Abstract

We develop an empirical framework to measure the impact of firms' endogenous quality choices on their production. Our approach provides unbiased estimates of productivity, whereas traditional methods would misattribute lower-quality care to higher productivity. In our application, we find a significant quality-quantity tradeoff for dialysis treatments: facilities may treat 1 percent more patients by allowing their expected infection rate to increase by 0.8 percentage points (roughly 6 percent), holding inputs and patient characteristics fixed. We also find (i) extensive quality-adjusted productivity dispersion across providers, (ii) better outcomes among non-profit entities, and (iii) comparatively little effect from competition.

JEL: D24, I1, L2
Keywords: productivity; quality variation; health care

*We thank Matt Grennan, David Rivers, Mark Roberts, Frederic Warzynski, and participants at the 2012 International Industrial Organization Conference for helpful comments.

†The Pennsylvania State University, Department of Economics, paul.grieco@psu.edu

‡University of Rochester, Simon Graduate School of Business, ryan.mcdevitt@simon.rochester.edu

1 Introduction

Rising healthcare expenditures have prompted spending reforms such as Medicare's prospective payment system, which ties reimbursements to a fixed amount per service rather than to a provider's actual costs. While these reforms aim to limit wasteful healthcare expenses, they potentially lead to an unintended consequence: providers may reduce their quality of care to minimize costs and treat more patients. Determining whether providers do, in fact, trade off quantity and quality faces two challenges, however. First, measuring inputs, outputs, and quality in health care is often not straightforward. Second, providers' endogenous choices with respect to inputs and quality may bias estimates of productivity. To overcome these challenges, we develop a novel method to estimate a healthcare provision function that incorporates both a quality-quantity tradeoff and unobserved firm productivity. We then apply this technique to a tractable yet representative healthcare setting well-suited for our approach, outpatient dialysis treatments. From our analysis, we find that providers reduce costs by sacrificing quality, as allowing a one standard deviation increase in a center's infection rate decreases its costs by the equivalent of three full-time employees.

We focus specifically on dialysis treatments — a process that cleans the blood of patients with end-stage renal disease (ESRD), or kidney failure — for two primary reasons. First, payments to dialysis facilities comprise a substantial portion of Medicare's expenditures each year, making it an important area for policy analysis in its own right. Second, several features of the dialysis industry make it well suited for a study of quality provision in health care. For instance, dialysis is predominately conducted in standalone facilities, which creates a well-defined market for our analysis. In addition, payments for treatment are largely uniform due to Medicare's prospective payment system and do not depend on quality, making it possible to isolate the effects of quality provision from price discrimination. Moreover, dialysis treatments follow a straightforward process related to stations and staff, which allows us to closely approximate a facility's production function. Facilities also choose input levels (i.e., staffing) and have observable differences in production (i.e., patient loads), which allows us to identify the relationship between inputs and outputs. Finally, facilities have observable differences in outcomes that relate directly to the quality of care they provide (e.g., infection and death rates), which allows us to connect a firm's productivity to its treatment quality, the primary aim of our research.

To uncover the cost of providing higher-quality care, we build on the structural methods

for estimating firm-level production functions first proposed by Olley & Pakes (1996), and later extended by Levinsohn & Petrin (2003), Ackerberg et al. (2006), Gandhi et al. (2009), and others. Conceptually, we adapt these methods to incorporate a "quality-choice" stage that comes after a firm's choices of labor and capital. That is, after acquiring capital and training workers, a manager observes his center's expected level of productivity and chooses the level of quality to provide by, for example, stipulating guidelines for the length of treatment, the cleanliness of equipment, and the degree of oversight. Because we do not directly observe firms' quality choices, however, we use observable measures of patient outcomes as a proxy for what those choices must have been — if high-quality care is more likely to result in better health outcomes, those outcomes are valid proxies for quality choices. And because we have multiple measures of health outcomes (in our case, a center's septic infection and mortality rate), we can use an instrumental variable approach to recover the impact of quality choices on output.

We use our results to investigate why dialysis centers provide different levels of treatment quality. Due to Medicare's prospective payment system, dialysis centers have an incentive to minimize the costs of treating patients, which may include providing low-quality care. Counteracting this incentive, however, are plausible motivations for providing high-quality treatments: centers must report quality statistics to Medicare which are then made public, and face intermittent inspections by state regulators (Ramanarayanan & Snyder 2011). In addition, patients have a choice over their dialysis providers, potentially leading centers to compete for patients by providing higher-quality care (Dai 2012). Finally, non-profit centers may have objectives for providing high-quality care unrelated to maximizing profits (Sloan 2000).

From our analysis, we find a substantial quality-quantity tradeoff for dialysis treatments: a center can increase its patient load by 1 percent by allowing a 0.8 percentage point higher septic infection rate, holding input levels and productivity constant. In addition, our approach allows us to recover estimates of total factor productivity for each firm that properly account for endogenous quality choices. We find substantial productivity dispersion in the industry that is not explained by differences in treatment quality. Finally, we investigate the determinants of quality in the industry and find that for-profit dialysis centers provide significantly worse care, with an infection rate 1.5 percentage points (roughly 12 percent) higher than their non-profit counterparts. At the same time, local competition does not appear to improve treatment quality. Overall, our results provide evidence that profit-based incentives to reduce costs in dialysis treatments may lead to lower-quality care, and that competition has a limited impact

on quality.

In addition to providing relevant policy analysis, this paper also contributes to the growing literature in empirical industrial organization on the estimation of production functions. These methods have a long history in economics, with many well-known econometric issues related to selection and simultaneity bias receiving considerable attention.[1] In light of this, recent work has developed structural techniques that use firms' observed input decisions to control for unobserved productivity shocks and overcome endogeneity problems.[2] We extend these methods to incorporate observable measures of output quality into the production function, which is necessary for healthcare applications. To our knowledge, we are the first to apply these methods to a healthcare setting with the goal of measuring a quality-quantity tradeoff.[3]

Firms' quality choices require direct consideration within health care because failing to control for differences in treatment quality may bias estimates of productivity. As an example, consider the econometrician who observes factory A producing twice as many widgets as factory B using exactly the same measurable inputs, such as labor hours. Accounting measures of productivity, such as simple labor-to-output ratios, would suggest that factory A is twice as productive. If widgets from factory A fail at a higher rate because its employees devote less attention to quality control, however, then the true difference in productivity is less than the factor of two suggested by a simple quantity measure.

Previous work assumes away differences in output quality by arguing that higher-quality products (i.e., those with fewer defects) command higher prices, which will then be captured by revenue-denominated output measures. The mapping of price to quality has many shortcomings, however, as prices reflect more than just quality (Griliches & Mairesse 1995). As noted in De Loecker (2012), for instance, exchange rates and market power influence prices in ways that may obscure quality differences. To address this concern, De Loecker uses quantity-denominated measures for output rather than revenue-denominated measures deflated by a price index, leaving aside issues of quality dispersion.[4] While indirect controls for quality are suitable for some settings, output quality is of first-order importance in many applications and estimation techniques that do not properly incorporate quality variation across firms result in misleading

[1] See Syverson (2011) for a recent review.

[2] See, for example, Olley & Pakes (1996), Ackerberg et al. (2006), and Levinsohn & Petrin (2003).

[3] Romley & Goldman (2011) consider quality choices among hospitals using a revealed-preference approach rather than outcome-based quality measures. Lee et al. (2012) use a structural approach to measure the impact of healthcare IT on hospital productivity, but do not consider output quality.

[4] A separate line of research, such as Fox & Smeets (2011), uses the quality of inputs (e.g., the education levels of employees) to explain productivity dispersion.

inference. Moreover, standard techniques to control for quality — such as price variation — may not adequately resolve the issue, particularly in healthcare settings such as dialysis where prices are set by Medicare rather than individual facilities and are unrelated to the quality of care provided (in fact, they are set before the quality of care is even chosen). These features thus serves as the motivation for our approach.

The remainder of our paper continues in the following section with a description of the outpatient dialysis industry and our data sources. Section 3 outlines our methods for estimating a production function in the presence of an endogenous quality choice. Section 4 presents our estimation results. Finally, Section 5 concludes with a discussion of our findings' implications.

2 Empirical Setting and Data Description

The demand for dialysis treatments comes from patients afflicted by end-stage renal disease (ESRD), a chronic condition characterized by functional kidney failure that results in death if not treated properly. Patients with ESRD effectively have only two treatment options, a kidney transplant or dialysis. Because of the long wait-list for transplants, however, nearly all ESRD patients at some point must undergo dialysis, a process which cleans the blood of waste and excess fluids. Patients can receive different dialysis modalities, with hemodialysis, a method that circulates a patient's blood through a filtering device before returning it to the body, constituting 90.4 percent of treatments (Center for Medicare and Medicaid Services).

Patients receiving dialysis in the United States primarily do so at free-standing dialysis facilities, which collectively comprise over 90 percent of the market (USRDS 2010).[5] Medicare's ESRD program, instituted by an act of Congress in 1973, covers the majority of these patients; notably, all patients with ESRD become eligible for Medicare coverage, regardless of age, and the program now encompasses over 400,000 individuals. Today, Medicare spends more than $20 billion a year on dialysis care — about $77,000 per patient annually — which constitutes more than six percent of all Medicare spending despite affecting fewer than one percent of Medicare patients (ProPublica 2011). Beginning in 1983, Medicare has paid dialysis providers a fixed, prospective payment — the "composite rate" — for each outpatient treatment delivered, up to a maximum of three sessions per week per patient. Initially, the payment rate did not adjust for quality, length of treatment, dialysis dose, or patient characteristics, though Medicare began to

[5]Other options for receiving dialysis include hospital emergency rooms and in-home treatments.

adjust payments based on patient characteristics in 2005.

Dialysis treatments require constant supervision by trained medical professionals, as patients must remain connected to a dialysis machine for 2-5 hours to filter impurities and remove excess fluid from their blood. Prior to treatment, staff connect the machine to a patient by inserting two lines into a vascular access and asses his condition. During treatment, staff must continually monitor patients to evaluate conditions (e.g., blood pressure) and to treat symptoms that arise (e.g., hypotension). Following treatment, staff disconnect a patient from the dialysis machine and assess his condition a final time before discharge. As a result of this hands-on care, the cost per patient treated necessarily increases with the average duration of treatment. Labor costs, which consist largely of nurses and technicians' wages, reflect this, accounting for approximately 70-75 percent of a facility's total variable costs (Ford & Kaserman 2000).

Centers employ different types of labor, with registered nurses (RNs) constituting the majority of staff. Technicians, who have less-extensive training than RNs, also treat patients but can do so with only a high-school diploma and in-house training (though they must eventually pass a state or national certification test). Notably, centers cannot quickly react to changes in productivity by hiring more workers due to training and certification requirements. Centers also must have board-certified physicians as medical directors, but often have no physician on site. Medicare does not mandate a specific staffing ratio for dialysis centers, although some states do.

Another significant decision for dialysis facilities is the number of stations to utilize. Centers vary widely in terms of size, ranging from 1 to 80 stations. Based on industry reports, a typical dialysis station costs $16,000 and has a useful life of approximately seven years (Imerman & Otto 2004).

In addition to labor and capital decisions, firms also decide how much effort to put towards providing high-quality care. For example, dialysis sessions require up to one hour of preparation and cleaning, which can be shortened according to a manager's discretion and can directly affect treatment outcomes. Importantly, patients undergoing dialysis face a high risk of septic infection due to the exposure of their blood during treatment, with the risk depending on the cleanliness of the dialysis center. The center likely has considerable control over its targeted infection rate, as health professionals who follow straightforward procedures can effectively drive the infection rate to zero (Pronovost et al. 2006). The decision to do so, however, comes with the tradeoff of treating fewer patients.

Because a facility's payments per treatment do not vary by the duration of treatment under

Medicare's prospective payment system, a facility's profit per treatment decreases as treatment times — and, hence, labor costs — increase. At the same time, the effectiveness and safety of dialysis increases with its duration; for instance, longer treatment cycles have been linked to lower mortality rates (ProPublica 2011). Centers thus face a tradeoff between increasing treatment quality and decreasing costs.[6] And though the costs of providing high-quality care are relatively clear, the benefits for dialysis centers are less straightforward. First, demand-side incentives appear weak because dialysis provides life-sustaining functions for patients, making their demand for treatments inelastic. Second, patients typically have few dialysis centers to choose from in any given market — the mean market share across the United States is 0.457 — and, since ESRD immobilizes those affected by it, travel costs limit market choice. Finally, as discussed above, Medicare's payment system provides no direct financial incentive for providing high-quality care.

Firms may, however, still have several possible incentives for delivering high-quality care. For instance, when a facility does face competition for patients, providing low-quality care may lead its patients to defect to other facilities that provide relatively better care and have excess capacity. Moreover, a facility that provides inadequate treatment may face increased regulatory scrutiny that further drives patients to competitors or results in decertification (the degree of regulatory scrutiny varies from state to state). Finally, some centers, particularly non-profit entities, may have motives to provide high-quality care unrelated to profitability.

Data Sources We use several sources of data for our analysis. Our primary dataset comes from the Centers for Medicare and Medicaid Services which contracts with the University of Michigan's Kidney Epidemiology and Cost Center to compile customized reports for each dialysis facility in the country. In December 2010, ProPublica, a non-profit organization dedicated to investigative journalism, obtained these reports under the Freedom of Information Act and posted them online. We systematically downloaded all individual reports covering 2004 — 2008 and constructed a usable dataset, which to our knowledge is the first time it has been used for research purposes.[7] The data include detailed center-level information on aggregated patient (e.g., age, gender, co-morbid conditions, etc.) and facility characteristics (e.g., number of stations

[6]Critics allege that facilities may sacrifice quality of care in pursuit of efficiency, turning over three to four shifts of patients a day. And while policy makers contend that technicians should not monitor more than four patients at once, patient-to-staff ratios exceed this guideline in many facilities. At the extreme, inspection reports allege that some clinics have allowed patients to soil themselves rather than interrupt dialysis (ProPublica 2011).

[7]Others, such as Ramanarayanan & Snyder (2011), use a more limited form of the same data.

Table 1: Summary Statistics.

Variable	Mean	St. Dev.
Patient Years	50.856	31.913
FTE Staff	13.496	7.933
Net Hiring	1.064	0.552
Zero Net Hiring	0.127	0.333
Stations	18.612	7.877
Septic Infection Rate	12.504	6.399
Death Rate Ratio	1.041	0.405
Number of Firms	4,270	
Number of Firm-Years	18,295	

and nurses, years in operation, etc.).

Table 1 presents selected summary statistics from the data, and several variables deserve note. First, Medicare analyzes individual patient records and calculates the number of patient-years each dialysis center serves (e.g., a patient treated at a center for six months is accounted for as one half of a patient-year). We use this variable as our measure of output, as it provides an accurate record of dialysis provision that accounts for partial years of service due to death, transfers, transplants, newly diagnosed patients, and so forth. We also use the number of full-time equivalent (a weighted mix of full-time and part-time) employees at each center and the number of dialysis stations as our measures of labor and capital inputs. The average number of dialysis stations used by a center is 18, making the purchase of a new machine a significant investment; reflecting this, firms have zero net investment for 90 percent of the center-year observations in the data. In terms of hiring, centers, on average, increase their staff by the equivalent of one full-time employee each year, while 12.7 percent of centers have no net change in employment in a given year.

We use a center's hospitalization rate from septic (blood) infections as our primary measure of quality, which averages 12.5 percent per year, with a standard deviation of over 6 percent. In addition to the septic infection rate, we use the ratio of deaths to expected deaths as an alternative measure of quality.[8] Importantly, we can also control for aggregate patient characteristics at each center which influence productivity and quality, which we discuss at length in Section 3.3.

The competitive environment faced by dialysis centers is highly variable, as shown in Table 2.

[8]The center-level expected death rate is calculated by Medicare using individual patient characteristics.

Table 2: Distribution of Competitors in Health Service Areas.

Num Comp.	N	Freq.	Cum.
0	4,789	0.2618	0.2618
1	3,223	0.1762	0.4379
2	1,828	0.0999	0.5379
3+	8,455	0.4621	1.0000

Following the healthcare literature, we use hospital service areas (HSA) as our market boundaries for dialysis centers. While roughly 20 percent of dialysis centers are monopolies within their HSA, the average number of centers in an area is 8.1; in addition, the mean patient-weighted market share across centers within an HSA is 0.45.

3 Measuring the Quality-Quantity Tradeoff in Dialysis

To measure the relationship between productivity and treatment quality, we propose and estimate a structural model of dialysis provision. In doing so, we account for both the standard endogeneity problems that arise when estimating production functions and the additional problem introduced by a firm's endogenous choice of treatment quality. The complication related to endogenous quality decisions stems from the unobserved (to the econometrician) choice made by firms that receive positive shocks to productivity: they may choose either to treat more patients, or to treat current patients more intensively. If highly productive firms elect to provide higher-quality care for their patients, naïve estimates of the quality-quantity tradeoff will be biased, leading us to underestimate the true cost of quality.

To control for this potential source of bias, we extend the work of Olley & Pakes (1996) and Ackerberg et al. (2006) by incorporating firms' endogenous quality targets. Because we only observe noisy measures of quality in our data, however, we also control for measurement error in quality choices, proxied for by firm-level hospitalization rates for septic infection in our application. Specifically, the attenuation bias introduced by measurement error in quality choices would cause us to underestimate the magnitude of the quality-quantity tradeoff, which we correct for using an IV approach.

3.1 The Production Technology

Following the production function literature, we assume that the technology for dialysis treatments follows a Cobb-Douglas form:

$$Y_{it} = A_{it}(q_{it}, \omega_{it}) K_{it}^{\beta_k} L_{it}^{\beta_\ell}, \tag{1}$$

where Y_{it} is the number of dialysis treatments provided by center i in period t; the capital input, K_{it}, is the number of dialysis stations in center i; the labor input, L_{it}, is the full-time equivalent workforce at the center; and $A_{it}(q_{it}, \omega_{it})$ is a Hicks-neutral technology shifter that depends on a scalar reflecting the effort a center puts towards quality, q_{it}, and a scalar representing the firm's total factor productivity, ω_{it}, which the firm observes.

Using a scalar variable to represent productivity is a common restriction in the literature, and we model a firm's quality target as a unidimensional variable as well. While clearly restrictive, this simplification allows us to connect a firm's quality target to observable outcomes in a direct manner. By increasing the effort it puts towards quality, the firm incurs additional costs but improves the likelihood of positive treatment outcomes. For instance, increasing q_{it} leads to a decrease in the center's expected infection and death rate in our application. At the same time, q_{it} does not capture all center-level characteristics that a patient might associate with "quality." For example, a center could provide nicer amenities that do not affect health outcomes, such as providing televisions for patients during dialysis, which are not captured as quality improvements within our model.

The technology shifter, $A(\cdot)$, has the following functional form,

$$A_{it}(q_{it}, \omega_{it}) = e^{\alpha_0 + \alpha_q q_{it} + \omega_{it} + \epsilon_{it}}, \tag{2}$$

where ϵ_{it} is an unanticipated productivity shock that is uncorrelated with all other variables (it may also account for measurement error). The parameter α_q measures the magnitude of the quality-quantity tradeoff, and is presumed to be negative. As such, improving quality forces a center to reduce the number of patients it treats for a given level of inputs.

Taking the logarithm of (1) and letting lower case letters stand for the logarithm of upper

case letters, we arrive at the linear equation,

$$y_{it} = \alpha_0 + \alpha_q q_{it} + \beta_k k_{it} + \beta_\ell \ell_{it} + \omega_{it} + \epsilon_{it}. \tag{3}$$

Equation (3) makes apparent the well-known endogeneity problem associated with estimating production functions: because ω_{it} is observed by the firm but not the econometrician, it may be correlated with the firm's capital and labor choices. Our approach adds an additional endogeneity problem, as ω_{it} may also affect the firm's quality target. As a result, OLS estimates of (3) are inconsistent. Classical methods of correcting for endogeneity involve applying instruments for capital, labor, and quality, or assuming productivity is fixed over time (i.e., $\omega_{it} = \omega_i$) and using a fixed-effects estimator (Mundalk 1961). In practice, these approaches have had limited success. While input prices would seem to be appropriate instruments for capital and labor choices, they often have weak predictive power and can be difficult to obtain. A valid instrument for quality targets that is uncorrelated with unobserved productivity would be even more challenging to find. Furthermore, while the fixed-effects assumption is relatively easy to implement, it is quite strong and would not resolve the endogeneity problems if changes in productivity are responsible for changes in input (or, in our case, quality) choices.

To address these issues, Olley & Pakes (1996) propose an explicit structural approach to estimate the production process which uses observed firm decisions as proxies for unobserved productivity shocks, with the basic ideas behind this method extended further by Levinsohn & Petrin (2003) and Ackerberg et al. (2006).[9] In practice, the detailed timing assumptions required for the structural approach must be carefully evaluated to determine whether they fit the industry and data under consideration. As such, we discuss the timing assumptions of our model and relate them to the dialysis industry in the following subsection.

3.2 The Timing of Dialysis Center Decision Making

In their seminal paper, Olley & Pakes (1996) use capital investment as a proxy for unobserved productivity under the motivation that firms with greater productivity, *ceteris paribus*, will make larger investments. Given this intuition, differences in investments will provide a useful indication of differences in productivity. While natural for their setting of telecommunications equipment, this approach is not appropriate for dialysis centers because investment in new

[9] A second approach to production function estimation comes from the dynamic panel literature (e.g., Blundell & Bond 2000); Ackerberg et al. (2006) provides a comparison of these approaches.

stations is too infrequent: investment is zero for over 90 percent of the firm-year observations in the data. In light of this, we instead use firms' hiring decisions, which provide a natural proxy in our setting. Nurses and technicians employed by dialysis centers require training and credentialing, which introduces costs and time lags to hiring and layoff decisions. Therefore, we regard labor as a dynamic variable, which allows us to use a firm's (net) hiring decision to recover ω_{it}.[10] In contrast to labor choices, a firm can quickly adjust the quality of care it provides. For example, to improve quality, a manager could advise his center's staff to take extra precautions when treating patients, or to reduce quality by placing less emphasis on cleanliness and more on speed (Pronovost et al. 2006). While a center can dictate these policy changes more quickly than it can make can make hiring or investment changes, a lag still exists between quality decisions and actual implementation.

A firm's manager makes investment, hiring, and quality choices based on his center's capital stock, labor productivity, and a vector of other observable characteristics, x_{it}. Note that the components of x_{it} may affect the firm's policy function even though they do not affect production directly, and may include the extent of competition in the market, the firm's taste for quality via it's non-profit status, and other related variables. This leads to the timing assumptions of our model:

1. In period $t-1$, firms observe their productivity, $\omega_{i,t-1}$, and state, $x_{i,t-1}$; realize output $y_{i,t-1}$; and make investment, $i_{i,t-1}$, and hiring, $h_{i,t-1}$, decisions.[11] Newly hired workers (and newly invested capital) do not become available until period t, making the transitions for labor and capital:

$$k_{i,t} = k_{i,t-1} + i_{i,t-1} \qquad \ell_{i,t} = \ell_{i,t-1} + h_{i,t-1}.$$

2. At time $t-b$, which lies between periods $t-1$ and t, the firm discovers its new observable state, $x_{i,t}$ (for example, it observes whether additional firms will enter between periods $t-1$ and t), and observes its "interim" productivity, $\omega_{i,t-b}$. The firm also chooses its quality target, $q_{i,t}$.

3. At time t, the firm observes $\omega_{i,t}$, realizes production, $y_{i,t}$, and makes its hiring and investment decisions, $i_{i,t}$ and $h_{i,t}$.

[10] Note that this assumption conflicts with OP's conception of labor representing an immediately flexible input, though the distinction fits our setting.

[11] Strictly speaking, the investment decision may be made at or before time $t-1$.

In line with the literature, we assume productivity follows an exogenous Markov process between periods $t-1, t-b$, and t:

$$E[\omega_{it-b}|I_{it-1}] = E[\omega_{it-b}|\omega_{it-1}], \qquad E[\omega_{it}|I_{it-b}] = E[\omega_{it}|\omega_{it-b}],$$

where I_{it} represents firm i's information set at time t.

In this setting, unobserved productivity encompasses any factor that allows a center to treat more patients given its observable characteristics and quality target. For instance, a center's patients may follow treatment protocols more closely than others, which then frees the center either (i) to treat more patients because it devotes less time to dealing with complications that arise, or (ii) to spend extra time treating existing patients more intensively, which ultimately improves outcomes but does not appear in raw productivity measures, such as output-to-labor ratios.

3.3 Estimation

We use the model to estimate the underlying parameters of the production function and recover each firm's unobserved productivity in every period. We assume that firms behave optimally given the information they have at the time of their decision, making the firm's hiring decision and quality target functions of its current state. In addition to the production function variables, we assume that the firm's state includes a vector of center characteristics, x_{it}. That is, we denote the firm's hiring and quality policies as,

$$h_{it} = h(\omega_{it}, k_{it}, \ell_{it}, x_{it}), \qquad q_{it} = q(\omega_{it-b}, k_{it}, \ell_{it}, x_{it}),$$

where x_{it} is a state variable of other factors that affect firm policies but do not directly enter the production function.

Estimation proceeds in three steps. First, because we do not observe quality directly, we must find an appropriate proxy for quality based on center-level outcomes. Second, we specify the observed policy shifters, x_{it}, which will be included in the firm's hiring function. Finally, we adapt the standard two-stage estimation strategy to incorporate an endogenous quality choice with a noisy proxy.

3.3.1 Proxy for the Quality Target

Although treatment quality is not observed directly, the data contain information on patient outcomes that are correlated with it. In particular, we focus on the center's infection rate as an indicator of quality. This represents an imperfect measure, however, because variation in the infection rate may be due to differences in patient characteristics across centers rather than differences in centers' quality choices. To account for this, we control for center-level averages of several patient characteristics that affect infection rates. Specifically, we use the (negative) residual from a regression of infection rates on patient characteristics as our proxy for patient quality. This residual represents the variation in infection rates that remains after controlling for differences in the patient pool, and therefore serves as a proxy for the center's target for providing high-quality treatments.

We control for several patient characteristics, with summary statistics displayed in Table 3, that influence a center's infection rates beyond its quality decisions. Most notably, we include controls for patients' methods of vascular access. Patients receive dialysis through three main types of vascular access: arteriovenous (AV) fistula, AV graft, and venous catheter. A patient's vascular access method influences the likelihood of developing a blood infection, with an AV fistula being significantly less likely to form clots or become infected. Centers vary in the proportion of patients with an AV fistula, which ultimately may affect treatment outcomes. In addition to a patient's method of vascular access, other characteristics may directly affect treatment outcomes. Because centers vary in terms of their patients' characteristics, we also include controls for patients' average number of comorbid conditions, average duration of ESRD, and average hemoglobin levels,[12] as well as the patients' average age and proportion of the center's patients who are female.

We use the residual from a regression on septic infection rates as our primary proxy for quality. To control for measurement error, we then employ a second proxy as an instrument for quality targets, a center's actual death rate relative to expectations. Specifically, we use Medicare's estimates for each center's expected death rate which is based on individual patient characteristics (individual-level characteristics are not released to protect patient privacy). While Medicare uses the ratio of actual deaths to expected deaths as an indicator of center quality in its own reports, the death-rate ratio is subject to a large amount of noise and we do not use it as our primary measure of quality. Instead, we use this ratio as an instrumental

[12] Low hemoglobin levels are associated with anemia and pose health risks for dialysis patients.

Table 3: Patient Characteristics Summary Statistics.

Variable	Mean	St. Dev.
Avg. Patient Age	61.518	4.381
Pct. Female	45.798	8.333
Pct. AV Fistula	43.016	13.477
Avg. Comorbid Conditions	3.026	0.826
Avg. Duration of ESRD	4.089	0.953
Avg. Hemoglobin Level	11.882	0.332
Number of Firm-Years	18,295	

variable for quality targets.

3.3.2 Controlling for Policy Shifters

In order to invert the hiring function and recover each firm's productivity, we must explicitly control for all factors that affect hiring other than productivity. In our specification, we include the following sources of variation in x:

For-profit Status Centers differ in their ownership type, with roughly 87.7 percent operating as for-profit entities and the remainder as non-profit. A center's ownership structure may affect its polices related to hiring and treatment quality, and we therefore control for this distinction by including a dummy variable for the center's for-profit status in x_{it}.

Competition Because demand for dialysis centers is local, the extent of competition a center faces may affect its hiring and quality choices. For instance, centers in highly competitive markets may choose to improve quality or increase staff levels to attract patients. We include the level of competition each center faces in x_{it} in the form of dummy variables for having 0, 1, 2, or 3 or more competitors in an HSA. We assume that entry is exogenous and realized at the beginning of the period, so the firm observes its competitors when making its quality and hiring choices.

3.3.3 Two-Step Estimation

Under the assumption that a firm's hiring policy is monotonically increasing in productivity, we can invert the hiring function to recover a firm's productivity as a non-parametric function of

observables:[13]

$$\omega_{it} = h^{-1}(h_{it}, k_{it}, \ell_{it}, x_{it}). \tag{4}$$

Likewise, firms choose an optimal quality target given their information set at time $t-b$,

$$q_{it} = q(\omega_{it-b}, k_{it}, \ell_{it}, x_{it}). \tag{5}$$

Note that k_{it} and ℓ_{it} have been determined already by virtue of the investment and hiring decisions at time $t-1$, while we assume that the observable state, x_{it}, is revealed to the centers by the intermediate period $t-b$. Our timing assumptions imply that the quality policy is *not* collinear with the hiring policy due to the innovation in productivity between time $t-1$ and $t-b$.[14] Moreover, we do not place any invertibility assumptions on $q(\cdot)$.

Substituting (4) into (3), we arrive at our first-stage estimating equation,

$$\begin{aligned} y_{it} &= \alpha_0 + \alpha_q q_{it} + \beta_k k_{it} + \beta_\ell \ell_{it} + h^{-1}(h_{it}, k_{it}, \ell_{it}, x_{it}) + \epsilon_{it} \\ &= \alpha_q q_{it} + \Phi(h_{it}, k_{it}, \ell_{it}, x_{it}) + \epsilon_{it}, \end{aligned} \tag{6}$$

where

$$\Phi(h_{it}, k_{it}, \ell_{it}, x_{it}) = \alpha_0 + \beta_k k_{it} + \beta_\ell \ell_{it} + h^{-1}(h_{it}, k_{it}, \ell_{it}, x_{it}). \tag{7}$$

Due to invertibility requirements, we only have observations of (6) whenever hiring is non-zero.[15] Recall that ϵ_{it} is independent of all the observables in this equation by the assumption that it is not revealed to the firm when its hiring, investment, or quality decisions are made. If we had a perfect measure of quality, α_q could be consistently estimated from this equation using a kernel-based partially linear estimator (Robinson 1988). As discussed above, however, we do not directly observe quality targets, but instead observe only quality-related outcomes (e.g., infection rates, hospitalization rates, death rates, etc.). Because we observe multiple noisy measures of quality, we employ one as an instrument to consistently recover α_q from a

[13] There do appear to be some adjustment costs to hiring, as centers hire no workers in roughly 18 percent of center-years. We drop these observations since the hiring function will not be invertible in this range, in line with Olley & Pakes (1996).

[14] Ackerberg et al. (2006) emphasize the importance of this point.

[15] Because there are likely adjustment costs to hiring, h^{-1} is not well defined when hiring is zero (multiple productivity levels may lead to zero net hiring). We follow the productivity literature and drop observations of zero hiring when estimating the first stage.

sieve-based estimator using instrumental variables (Chen 2007). In practice, we use the septic infection residual discussed above as an error-ridden proxy for q_{it} and instrument it with the ratio of expected to actual deaths.

We recover the remaining parameters in a second stage. Note that, given any $\beta = (\beta_k, \beta_\ell)$, we can compute an estimate of unobserved productivity for each firm-year that has non-zero hiring,

$$\hat{\omega}_{it}(\beta) = \hat{\Phi}(h_{it}, k_{it}, \ell_{it}, x_{it}) - \beta_k k_{it} - \beta_\ell \ell_{it}.$$

Because ω_{it} follows a Markov process,

$$\omega_{it} = g(\omega_{it-1}) + \xi_{it}, \tag{8}$$

where g is a non-parametric function of ω_{it-1}, and ξ_{it} is a shock to productivity between time $t-1$ and t that is independent of the center's time-t information set.[16] Thus, for any given $\beta = (\beta_k, \beta_\ell)$, we can estimate $g(\cdot)$ using the estimating equation:[17]

$$y_{it} - \hat{\alpha}_q q_{it} - \beta_k k_{it} - \beta_\ell \ell_{it} = g(\hat{\omega}_{it-1}(\beta)) + \eta_{it}(\beta),$$

which follows from substituting the production function from (3) into the productivity innovation from (8), where $\hat{\alpha}_q$ is the consistent estimator of α_q recovered in the first stage. At the true value of β, $\eta(\beta) = \epsilon_{it} + \xi_{it}$, and so, by construction, $\eta_{it}(\beta)$ is uncorrelated with the time-t labor and capital variables. Therefore, β can be consistently estimated using the moment conditions,

$$E\begin{bmatrix} \eta_{it}(\beta) k_{it} \\ \eta_{it}(\beta) \ell_{it} \end{bmatrix} = 0. \tag{9}$$

We use (9) to estimate $\hat{\beta}$ via GMM, which can then be used to recover firm-level productivity estimates.

Standard errors are calculated using the block bootstrap, which accounts for statistical uncertainty in recovering the quality proxy, as well as both stages of the estimation process.

[16] Since we normalize the mean of ω_{it} to be zero, the constant term of g is a consistent estimator of α_0.

[17] We can estimate this equation using each observation that follows a non-zero hiring period. While it might seem more straightforward to recover $g(\cdot)$ by regressing $\hat{\omega}_{it}(\beta)$ on $\hat{\omega}_{it-1}(\beta)$, this would introduce a selection problem because we would only be able to use observations where hiring in period t itself was non-zero. We thank David Rivers for pointing this out to us.

4 Results

4.1 Production Function Estimates and the Quality-Quantity Tradeoff

In order to compare our structural method to OLS and fixed-effects (FE) procedures, we present the results from estimates of dialysis centers' production functions using each technique in Table 4. For the results relating to the structural model, we use a fifth-order polynomial with interactions to approximate $\Phi(\cdot)$ in the first stage, and a fifth-order polynomial to approximate $g(\cdot)$ in the second stage.

The first three columns in Table 4 present results from specifications that do not include the infection rate as a proxy for quality, with the estimates of β_k and β_ℓ differing substantially across the three estimation methods. Comparing our structural estimation in Column (III) to OLS in Column (I) and FE in Column (II) highlights several distinguishing features. First, OLS does not control for endogenous input choices, biasing the capital coefficient downwards and the labor coefficient upwards. This bias occurs because OLS relies on cross-sectional variation in stations to identify the labor and capital coefficients while ignoring the possibility of productivity differences across firms.

The FE procedure, in contrast, assumes productivity differences across firms remain constant over time, and estimates the capital and labor coefficients on the basis of year-to-year changes in center inputs. Using this method, both the capital and labor coefficients fall substantially for two primary reasons. First, relying on only year-to-year variation makes measurement error in both capital and labor inputs more of a concern. Because stations and employees remain fairly stable over time, attenuation bias from measurement error in the changes to hiring and investment biases these coefficients towards zero.[18] A second potential reason for the discrepancy between the OLS and FE approaches is that capital and labor differences in the cross section may proxy for unobserved, time-invariant center characteristics (e.g., center size) that the FE specification captures though the productivity term. That the OLS results indicate the production function has increasing returns to scale best illustrates this distinction, as we would expect that increasing the number of stations and staff within a center of constant size to exhibit decreasing returns to scale.

Finally, the third column presents results from estimates of the model presented in Section

[18] For example, if a new station was installed in June of 2002, it will first be reported in 2003, but the difference in the number of patients served in 2002 versus 2003 will underreport the impact of the new station that actually came online for the second half of 2002.

Table 4: Production Function Estimates.

	Without Quality Choice			With Quality Choice		
	OLS	FE	Model	OLS	FE	Model
Quality Effort, α_q				-0.0028	-0.0018	-0.0124
				(0.0007)	(0.0004)	(0.0042)
Capital, β_k	0.4628	0.1798	0.5210	0.4607	0.1788	0.5134
	(0.0208)	(0.0513)	(0.0446)	(0.0209)	(0.0514)	(0.0468)
Labor, β_ℓ	0.6709	0.1846	0.2527	0.6723	0.1855	0.2453
	(0.0149)	(0.0118)	(0.0304)	(0.0149)	(0.0119)	(0.0319)

3 with the added restriction that $\alpha_q = 0$. This specification employs a Markov process for productivity and uses both cross-section and time-series variation to identify the parameters, and utilizes firms' hiring choices to identify unobserved productivity. These results exhibit decreasing returns to scale with respect to stations and staff, as expected. In addition, they indicate that the impact of additional machines is roughly twice that of increasing the number of staff, which seems natural given the production technology for dialysis procedures. While increasing the number of employees may allow a firm to treat more patients by speeding up the transition of a dialysis station from one patient to another, the number of patients being treated by the center at any given time is necessarily bounded by the number of stations in the center.

We next turn to the primary focus of the paper, estimating the quality-quantity tradeoff for dialysis centers, α_q. The final three columns of Table 4 present results from specifications that control for treatment quality using OLS, FE, and our structural model. All three specifications provide evidence of a statistically significant quantity-quality tradeoff. However, the magnitude of the effect is much larger when using the structural model than with either the OLS or FE methods. The coefficient of -0.0124 from the structural model indicates that, holding inputs fixed, a firm that improves its quality enough that its expected infection rate falls by 1 percentage point will need to reduce overall patient hours by 1.24 percent, holding inputs fixed. Alternatively, the center could raise quality by the same amount and maintain the current level of output by increasing the size of its staff by roughly 5.1 percent. Given that the average center employs approximately 11 full-time equivalent nurses, this roughly equates to hiring one additional part-time worker. To lower its expected infection rate by a full standard deviation (6.3 percentage points) would cost the equivalent of roughly three full-time workers.

The smaller impact of quality on output in the OLS and FE specifications likely stems

Table 5: Robustness Checks.

	I	II	II	IV
Quality Effort, α_q	-0.0124	-0.0106	-0.0101	-0.0121
	(0.0042)	(0.0036)	(0.0042)	(0.0009)
Capital, β_k	0.5134	0.5077	0.4381	0.5136
	(0.0468)	(0.0474)	(0.0553)	(0.0466)
Labor, β_ℓ	0.2453	0.2448	0.1989	0.2455
	(0.0319)	(0.0313)	(0.0169)	(0.0318)
Control for Patient Characteristics	YES	NO	YES	YES
Control for Market Characteristics	YES	YES	NO	YES
Instrument for Quality	YES	YES	YES	NO

from endogeneity bias. We would expect, and in fact verify below in Table 8, that quality provision is positively associated with productivity. Since the OLS specification does not control for differences in productivity, we would expect its estimate of α_q to be biased upward (i.e., less negative). While the FE approach controls for time-invariant productivity levels, if firms' changes in quality targets are positively correlated with changes in their productivity, the FE estimate of α_q will also be biased upwards. This effect, coupled with the effects of attenuation bias already discussed above, bias the estimates of the quality-quantity tradeoff towards zero.

In Table 5, we consider several robustness checks of the baseline results, which are repeated in the first column. The second column drops controls for patient characteristics and instead simply uses the infection rate itself as a proxy for quality targets. The third column drops the center characteristics of for-profit status and competition from the hiring function. Finally, the fourth column does not instrument for the quality proxy but instead simply uses OLS to estimate the first stage. In all cases, the effect of quality declines slightly, though the finding of a significant quality-quantity tradeoff remains robust to various model specifications.

4.2 Productivity Dispersion, Growth, and Persistance

Having estimated the firm-level production function, we are able to recover center-year (log) productivity from

$$\hat{\omega}_{it} = y_{it} - \hat{\alpha}_0 - \hat{\alpha}_q q_{it} - \hat{\beta}_k k_{it} - \hat{\beta}_\ell \ell_{it}.$$

This allows us to analyze the dispersion, growth, and persistence of productivity within the dialysis industry. Moreover, we are able to estimate the importance of productivity for firms'

quality decisions.

To asses the extent of productivity dispersion, we first calculate the proportion of the variance in output that is explained by the production function outside of productivity differences:

$$R^2 = 1 - \frac{V(\hat{\omega}_{it})}{V(y_{it})}.$$

Our results indicate that the amount of productivity dispersion in the dialysis industry is substantial, with $R^2 = 0.489$, meaning that about half of the variation in output is attributable to productivity differences across firms, not input or quality differences. For a basis of comparison, Fox & Smeets (2011) report R^2 statistics for service industries ranging from 0.438 (Accounting) to 0.739 (Computer Activities).

We can then use these productivity estimates to measure productivity growth and persistence within the dialysis industry, as reported in Table 6. Overall, average productivity for the industry is roughly constant over the sample, with a slight drop in 2008. Again we find significant productivity dispersion across the industry. The inter-quartile range indicates that a firm in the 75th percentile of productivity is over 50 percent more productive than one in the 25th percentile.

On average, productivity growth at the firm level is extensive, ranging between 4 and 7 percent per year; at the same time, we observe a large degree of variation in productivity growth within the sample. The contrast of large firm-level productivity growth with slow industry-wide productivity growth suggests that firms enter at a lower level of productivity than their more-established peers.[19] Despite the high average growth rates, there is substantial dispersion in growth rates across centers, suggesting that year-to-year productivity shocks have substantial impact on centers' output. These shocks could result from high staff turnover, changes in patient characteristics, or other factors that affect productivity. We also find that productivity is persistent within a firm across years, as shown by a correlation in log productivity of approximately 0.8 for the entire sample.

We further explore the trends in productivity across firms in Table 7. Here, we stratify centers by age, determined by the year in which they first appear in the sample.[20] The average productivity of firms increases substantially with age, while the dispersion in productivity falls

[19]The decline in the number firms in 2008 is due to incomplete reporting of centers' staffing levels rather than actual closures, which are rare in this industry. For the purposes of estimation, we assume these data are missing at random.

[20]Centers appearing in 2004 are assumed to be four or more years old.

Table 6: (Log) Productivity, Productivity Growth, and Persistence

Year	N	Mean Level	St. Dev. Level	IQR Level	Mean Gain	St. Dev. Gain	Corr($\omega_{it}, \omega_{i,t-1}$)
2004	3,360	1.6646	0.5350	0.5043			
2005	3,563	1.6667	0.5279	0.5128	0.0682	0.3081	0.8191
2006	3,733	1.6783	0.4879	0.5151	0.0511	0.3113	0.7712
2007	3,885	1.6615	0.4979	0.5121	0.0385	0.2843	0.8130
2008	3,754	1.6177	0.5283	0.5253	0.0406	0.3020	0.7935
Total	18,295	1.6575	0.5155	0.5151	0.0492	0.3015	0.7995

Notes: Mean level is the average log productivity of all centers active in year t. Mean gain is the average change in log productivity of centers active in years t and $t-1$.

Table 7: (Log) Productivity Growth of New Firms

Age (years)	N	Mean Level	St. Dev. Level	IQR Level	Mean Gain	St. Dev. Gain	Corr($\omega_{it}, \omega_{i,t-1}$)
0	910	0.6858	1.0115	1.3557			
1	668	1.4528	0.4673	0.6018	0.6652	0.7363	0.6947
2	447	1.5633	0.4098	0.5309	0.1012	0.1985	0.9000
3	232	1.6041	0.3973	0.4970	0.0450	0.1968	0.8812
4+	16,038	1.7246	0.4111	0.4644	0.0147	0.2199	0.8475
Total	18,295	1.6575	0.5155	0.5151	0.0492	0.3015	0.7995

with each age group. Note, however, that the increase in productivity from age 0 to 1 is at least partially due to centers only operating for a portion of their initial year, and the results for productivity growth in years 1-3 indicate a fast but declining rate of productivity growth over the initial years of a center's existence. In contrast, the average productivity growth rate of 1.5 percent indicates only modest growth for established firms. Overall, it appears that new firms enter with productivity levels well below the industry average, but then experience strong growth to "catch up" to established firms. This, coupled with the slow growth of established firms themselves, results in relatively stagnant productivity growth for the industry as a whole.

4.3 The Determinants of Quality

The quality-quantity tradeoff estimated in Section 4.1 underscores the costs of providing high-quality care. To understand the mechanisms that lead firms to incur such costs, we next turn

Table 8: Linear Quality Regressions.

	I	II	III	IV
Productivity	0.2544	0.2885	0.1879	0.2278
	(0.7824)	(0.7761)	(0.7885)	(0.7823)
Capital	-0.5047	-0.3013	-0.3412	-0.1584
	(0.2509)	(0.2469)	(0.2463)	(0.2425)
Labor	-0.1796	-0.3313	-0.1433	-0.2987
	(0.4592)	(0.4551)	(0.4660)	(0.4627)
For Profit		-1.5754		-1.5428
		(0.2068)		(0.2069)
Monopolist			0.6374	0.5732
			(0.2053)	(0.2041)
Duopolist			-0.3147	-0.3601
			(0.1960)	(0.1928)
Triopolist			-0.4300	-0.4234
			(0.2322)	(0.2296)
Constant	1.4519	2.5725	0.9403	2.1145
	(0.6993)	(0.7020)	(0.6687)	(0.6796)

Note: Productivity, Labor, and Capital are expressed in logs. Dependent variable is the negative residual of a regression of centers' infection rate on average patient characteristics.

to the policy function for quality itself. From this, we can assess whether non-profit firms, which may have objectives beyond maximizing profits, tend to offer higher-quality treatments. Alternatively, we can determine whether competition leads firms to improve care.

To do so, recall that our model implies that the firm's quality policy is a nonparametric function of the tuple $(k_{it}, \ell_{it}, x_{it}, \omega_{it-b})$. To summarize the relationship between quality and its determinants, we regress our proxy for quality on $(k_{it}, \ell_{it}, x_{it})$ and our recovered estimate of productivity, $\hat{\omega}_{it}$. While this regression suffers from specification and measurement error (both q_{it} and $\hat{\omega}_{it}$ are contaminated with measurement error), it remains indicative of centers' quality policies. Again, we compute standard errors using the block bootstrap, which incorporates statistical sampling uncertainty in estimates of quality and productivity, and controls for firm-level serial correlation.

We present several specifications of this quality regression in Table 8. Across all specifications, quality is increasing in productivity, although the relationship is not statistically significant. The lack of precision of this estimate is perhaps not surprising given that both quality

Table 9: Partially Linear Quality Regressions.

	I	II	III	IV	V
For Profit	-1.5603		-1.5390	-1.5444	
	(0.2021)		(0.2030)	(0.2111)	
Monopolist		0.5390	0.4824		0.4725
		(0.2211)	(0.2196)		(0.2222)
Duopolist		-0.2474	-0.2977		-0.2926
		(0.1876)	(0.1843)		(0.1855)
Triopolist		-0.4701	-0.4678		-0.4431
		(0.2257)	(0.2234)		(0.2224)
Nonparametric Control for:					
Productivity	Yes	Yes	Yes	Yes	Yes
Capital	Yes	Yes	Yes	Yes	Yes
Labor	Yes	Yes	Yes	Yes	Yes
For-Profit Status	No	No	No	No	Yes
Competition	No	No	No	Yes	No

and productivity are measured with error. Moreover, attenuation bias drives our estimate of the impact of productivity towards zero. Nonetheless, the results provide further evidence that productivity and quality are positively correlated. The results also indicate that quality is decreasing in firm size, measured either in terms of employment or capital stock.

Columns II and IV show that for-profit firms provide significantly worse care than non-profit firms, with the expected infection rate more than 1.5 percentage points (over 10 percent) higher at for-profits. This estimate is statistically significant, and provides strong evidence that firms respond to profit-based incentives by choosing to offer a lower quality of service.

Columns III and IV allow us to examine the impact of competition on quality, where the base category is centers with three or more competitors in their market (defined as an HSA). The results show no clear pattern between competition and firms' quality choices. While monopolists tend to provide *higher* quality than firms in more competitive markets, duopolists and triopolists offer weakly lower quality, which runs counter to conventional intuition. Overall, competition does not provide a strong incentive for firms to increase their quality of service.

As a robustness check on the linear specification used in Table 8, we use a partially linear model to compute the effects of for-profit status and competition while controlling non-parametrically for the other determinants of quality. The results, presented in Table 9, show that our earlier conclusions are largely robust to using this more flexible specification. Again, we see that for-profit status strongly affects quality choices, while the impact of competition is

much weaker and non-monotonic.

5 Conclusion

By estimating center-level production functions that incorporate endogenous quality choices, we find evidence that dialysis centers face a tradeoff between serving more patients and providing higher quality care. This result suggests that policies aimed at increasing efficiency may inadvertently affect health outcomes. Although we find considerable dispersion in productivity across firms, these results imply that incentives to reduce costs may lead to lower-quality care, not greater efficiency. Similarly, our results on non-profit centers also provide evidence that firms react to cost incentives by adjusting the quality of their treatments. Non-profit centers, which have less incentive to reduce costs, provide higher-quality care to patients than their for-profit counterparts, on average.

We find little evidence that market forces discipline centers to provide high-quality care. While competition might be expected to provide a demand-side incentive for improving quality, we find that firms in more-competitive markets are not more likely to offer better care than monopolists. Disentangling the potential explanations for this result lies beyond this paper, though the inelastic demand for dialysis treatments, the dominance of two for-profit chains, and the weak incentives imposed by Medicare all likely contribute to this outcome.

Because dialysis treatments comprise a large — and growing — cost for Medicare, controlling the cost of dialysis provision will likely concern policy makers for the foreseeable future. Our work informs policy by showing that, while productivity dispersion is extensive within the industry, cost-cutting initiatives may simply reduce the quality of care provided rather than promote efficiency. More importantly, because dialysis resembles other healthcare settings, these findings illustrate the challenges of introducing policies intended to minimize costs while maintaining high standards of care.

References

Ackerberg, D. A., Caves, K. & Frazer, G. (2006), Structural identification of production function. UCLA, Deloitte and Touche, Rotman School of Management.

Blundell, R. & Bond, S. (2000), 'Gmm estimation with persistent panel data: an application to production functions', *Econometric Reviews* **19**, 321–340.

Chen, X. (2007), Large sample sieve estimation of semi-nonparametric models, *in* J. J. Heckman & E. E. Leamer, eds, 'Handbook of Econometrics', Vol. 6B, Elsevier Science B.V., Amsterdam, chapter 74, pp. 5549–5632.

Dai, M. (2012), Specialization and differentiation in outpatient dialysis. Drexel Lebow.

De Loecker, J. (2012), 'Product differentiation, multi-product firms and estimating the impact of trade liberalization on productivity', *Econometrica* **forthcoming**.

Ford, J. M. & Kaserman, D. L. (2000), 'Ownership structure and the quality of medical care: evidence from the dialysis industry', *Journal of Economic Behavior & Organization* **43**, 279–293.

Fox, J. & Smeets, V. (2011), 'Does input quality drive measured productivity differences in firm productivity', *International Economic Review* **52**(4), 961–989.

Gandhi, A., Navarro, S. & Rivers, D. (2009), Identifying production function using restrictions from economic theory. University of Wisconsin-Madison.

Griliches, Z. & Mairesse, J. (1995), Production functions: The search for identification. NBER Working Paper 5067.

Imerman, M. & Otto, D. (2004), Preliminary Market and Cost Analysis of a Five-Station Hemodialysis Facility in Marengo, Iowa. Iowa State University.

Lee, J., McCullough, J. S. & Town, R. J. (2012), The impact of health information technology on hospital productivity. National Bureau of Economic Research.

Levinsohn, J. & Petrin, A. (2003), 'Estimating production functions using inputs to control for unobservables', *The Review of Economic Studies* **70**(2), pp. 317–341.
URL: *http://www.jstor.org/stable/3648636*

Mundalk, Y. (1961), 'Empirical production function free of management bias', *Journal of Farm Economics* **43**, 44–56.

Olley, G. S. & Pakes, A. (1996), 'The dynamics of productivity in the telecommunications equipment industry', *Econometrica* **64**(6), pp. 1263–1297.
URL: *http://www.jstor.org/stable/2171831*

Pronovost, P., Needham, D., Berenholtz, S., Sinopoli, D., Chu, H., Cosgrove, S., Sexton, B., Hyzy, R., Welsh, R., Roth, G. et al. (2006), 'An intervention to decrease catheter-related bloodstream infections in the icu', *New England Journal of Medicine* **355**(26), 2725–2732.

ProPublica (2011), 'Dialysis: High costs and hidden perils of a treatment guaranteed to all'.
URL: *http://www.propublica.org/series/dialysis*

Ramanarayanan, S. & Snyder, J. (2011), Reputations and firm performance: Evidence from the dialysis industry. UCLA Anderson.

Robinson, P. (1988), 'Root-n-consistent semiparametric regression', *Econometrica* **56**(4), 931–954.

Romley, J. A. & Goldman, D. P. (2011), 'How costly is hospital quality? a revealed-preference approach', *The Journal of Industrial Economics* **59**(4), 578–608.
URL: *http://dx.doi.org/10.1111/j.1467-6451.2011.00468.x*

Sloan, F. (2000), Not-for-profit ownership and hospital behavior, *in* A. Culyer & J. Newhouse, eds, 'Handbook of Health Economics', Vol. 1B, Elsevier Science B.V., Amsterdam.

Syverson, C. (2011), 'What determines productivity?', *Journal of Economic Literature* **49**(2), 326–65.

USRDS (2010), 2010 Annual Data Report, Technical report, United States Renal Data System, Minneapolis, MN.